

Detecting very long-lived gravitational-wave transients lasting hours to weeks

Eric Thrane,^{1,2, a} Vuk Mandic,³ and Nelson Christensen⁴

¹*LIGO Laboratory, California Institute of Technology, MS 100-36, Pasadena, CA, 91125, USA*

²*School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia*

³*School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA*

⁴*Physics and Astronomy, Carleton College, Northfield, Minnesota 55057, USA*

We explore the possibility of very long-lived gravitational-wave transients (and detector artifacts) lasting hours to weeks. Such very long signals are both interesting in their own right and as a potential source of systematic error in searches for persistent signals, e.g., from a stochastic gravitational-wave background. We review possible mechanisms for emission on these time scales and discuss computational challenges associated with their detection: namely, the substantial volume of data involved in a search for very long transients can require vast computer memory and processing time. These computational difficulties can be addressed through a form of data compression known as coarse-graining, in which information about narrow frequency bins is discarded in order to reduce the computational requirements of a search. Using data compression, we demonstrate an efficient radiometer (cross-correlation) algorithm for the detection of very long transients. In the process, we identify features of a very long transient search (related to the rotation of the Earth) that make it more complicated than a search for shorter transient signals. We implement suitable solutions.

I. INTRODUCTION

Previous work has explored the astrophysics of (and associated detection strategies for) long-lived gravitational-wave transients lasting ≈ 10 – 1000 s; see, e.g., [1–11] and references therein. Most long-transient models rely on the emission of gravitational waves from rotational instabilities in a protoneutron star or its accretion disk but other models are possible as well [1]. Initial LIGO/Virgo have performed a search for unmodeled long lived transients associated with long gamma ray bursts yielding upper limits on gravitational-wave fluence [12].

Despite this progress, there is still a largely unexplored region of parameter space of “very long-lived” transient signals lasting hours to weeks; (although, see [13]). As we discuss below, it may be possible for neutron stars to emit transient gravitational waves on these time scales. Moreover, exotic models allow for the possibility of a seemingly persistent signal to start or stop during an observing run [14], also potentially leading to very long transient signals. Thus, there is astrophysical motivation to carry out a search for very long lived gravitational wave transients. An efficient very long transient detection algorithm will have other useful applications: it can establish if an apparently persistent source, e.g., observed in a stochastic background search [15], exhibits variability in time, and it can be used to understand the behavior of detector artifacts on timescales of days to weeks.

Our goal here is to outline a computationally feasible search for very long-lived transients employing radiometry [16]. We review the principles of radiometry below in Section III, but for now, we note that radiometry relies on the cross-correlation of two or more detectors in order to detect gravitational waves as excess coherence. Ex-

cess coherence can be identified through various pattern recognition techniques. For demonstrative purposes, we here employ a “seedless clustering” algorithm [10, 11], discussed in greater detail in Section III, though, other pattern recognition techniques may be used as well.

In order to carry out such a search, it is necessary to first address computational challenges that appear, at first glance, to make the search daunting. Namely, difficulties arise due to the fact that only so much uncompressed spectrographic strain data can be stored at once in random access memory (RAM). Without compression, it is impossible to hold many hours of data in memory at once on a typical computer. In theory, data can be read in as needed, but this creates an input/output bottleneck. Also, in searches that rely on seedless clustering [10, 11], the computation time required to sum hours of uncompressed spectrographic data can become prohibitive.

The solution we propose is an extra pre-processing step in which data are compressed in order to facilitate computationally feasible analyses. Through a procedure known as coarse-graining (already used in searches for the stochastic background [17]), the volume of data can be reduced by orders of magnitude so that days of data can be represented with a spectrogram comparable to one used to represent minutes of uncompressed data. Coarse-graining necessarily discards information, but through this procedure, we may study much longer transient signals than otherwise possible.

The remainder of this paper is organized as follows. In Section II, we discuss the motivation for searches for very long transients. In Section III, we review key principles of radiometry needed for subsequent discussion. (Additional details are included in the Appendix.) In Section IV, we describe compression procedures for coarse-graining, and in Section V, we demonstrate the recovery of very long-lived signals using compressed data. In Section VI, we summarize our results and suggest some lines of future research.

^aElectronic address: eric.thrane@monash.edu

II. MOTIVATION

The scientific rationale for a search for very-long transients, spanning hours to weeks, was first explored in [13]. The authors of [13] review a number of somewhat speculative scenarios associated with neutron stars including gravitational-wave emission lasting days to months from non-axisymmetric Ekman flow following a glitch [18–20], Alfvén oscillations from giant magnetic flares (also lasting days to months) [21, 22], emission from free precession (with a damping time possibly lasting from weeks to years) [23–25], magnetic instabilities in newborn neutron stars (lasting days) [26], and gravitational-waves from r -modes [27, 28].

Generic rotational instabilities in newborn neutron stars, potentially powered by fallback accretion [8, 29], may persist on a timescale of hours [7]. Somewhat more speculatively, we note that observations of intermittent pulsars, which become quiescent on timescales of days (e.g., [30]), suggest that neutron star dynamics vary on the timescales considered here, motivating exploration of this region of parameter space. Similarly, variability in accretion on these timescales may affect gravitational-wave emission from accretion-supported mountains [31]. Finally, it is worthwhile to be prepared for a surprise: a very long lived transient signal from an unexpected source. Recent work proposing gravitational-wave emission from gravitationally bound axion clouds [14], potentially starting and stopping on the timescale of a few years, serves to illustrate this possibility. The method we propose below is also applicable to quasi-infinite signals turning on or off during an observing run as well as repeating sources with long-lasting emission periods.

An algorithm for identification of very long transients also provides a means of understanding the time dependence of apparently persistent signals, e.g., in a stochastic background search. In the event of an ostensibly persistent gravitational-wave detection, it is prudent to investigate if the temporal behavior of the signal is consistent with the assumed signal model. Matched filter searches for rotating neutron stars, for example, assume that the strain is constant over the duration of the measurement. An observation of non-trivial time-dependence could be evidence of new physics.

In addition to the astrophysical motivation, the algorithm provides useful information for detector characterization. Recent work has shown that geophysical effects such as global Schumann resonances can induce correlated noise in worldwide networks of gravitational-wave detectors in a way that can bias cross-correlation searches [32, 33]. Some geophysical phenomena (such as electromagnetic activity induced by solar flares) exhibit temporal fluctuations on very long timescales. Understanding the time dependence of excess cross correlation can help to rule out (or lend credence to) different explanations of an apparent signal in searches for persistent gravitational waves in both targeted radiometer searches [34] and searches for stochastic backgrounds [35].

In order to see how very long signals might escape detection without a dedicated search, we note simple scaling behavior. Radiometric signal-to-noise ratio (SNR, defined below; or see Eq. 3.11 of [1]) scales like

$$\text{SNR} \propto \zeta t_{\text{obs}}^{1/2}, \quad (1)$$

where $t_{\text{obs}}^{1/2}$ is the observation duration and ζ is the fraction of the observation during which the signal is present. To be concrete, let us consider an ≈ 11 hr-long signal with $\text{SNR}_{\text{tot}} = 40$. While this signal is exceptionally loud, it can be significantly diluted if we search for it assuming a persistent signal model. For example, if we integrate over the entirety of a 90 day science run, the signal-to-noise ratio will be diluted by a factor

$$\frac{(1)\sqrt{11 \text{ hr}}}{(11 \text{ hr}/90 \text{ days})\sqrt{90 \text{ days}}} = \sqrt{\frac{90 \text{ days}}{11 \text{ hr}}} = 14,$$

so that SNR_{tot} is reduced from 40 to ≈ 3 , which is not statistically significant given a few thousand independent frequency bins. It follows that a very long transient signal, loud enough to affect a radiometer or stochastic search, should produce an easily identifiable signal using an algorithm dedicated to detection of very long transient phenomena.

If we attempt to resolve the signal using 1 s segments, the signal-to-noise ratio is diluted by a factor of

$$\frac{(1)\sqrt{11 \text{ hr}}}{(1)\sqrt{1 \text{ s}}} \approx 200,$$

so that $\text{SNR} \approx 0.20$ in each segment. This back-of-the-envelope argument illustrates how even exceptionally loud very long signals can escape detection without a dedicated search.

III. RADIOMETRY

The basic idea of radiometry is to cross-correlate data from two or more detectors. By integrating over as much data as possible, it is possible to identify weak signals buried in noise. Drawing on previous results [1, 10, 11], we take as our starting point spectrograms of cross-correlated data. Details are provided in Appendix A. For the sake of simplicity, we focus here on a two-detector network, but note that the algorithm generalizes straightforwardly to arbitrarily many detectors [1]. The cross-correlated data can be represented as a complex-valued signal-to-noise ratio [11] (see Eq. A4 in Appendix A):

$$\mathbf{p}(t; f) = \kappa \tilde{s}_1^*(t; f) \tilde{s}_2(t; f) / \sqrt{P'_1(t; f) P'_2(t; f)}. \quad (2)$$

Here, $\tilde{s}_I(t; f)$ is the discrete Fourier transform of the strain series from detector I , $P'_I(t; f)$ is the auto-power spectrum calculated using neighboring data segments, and κ is a normalization factor; (see Eq. A5 in the Appendix). Each segment has duration δt . The choice of δt ,

which we examine in greater detail below, is determined both by astrophysical considerations and computational limitations. The argument t is the mid time of the data segment and f is frequency. Note that $\mathbf{p}(t; f)$ does not depend on the direction of the source $\hat{\Omega}$ because the direction of the source is encoded as a complex phase.

If the data consist of well-behaved noise, both the real and imaginary parts of \mathbf{p} are characterized by distributions with mean=0 and variance \approx 1. Gravitational-wave signals induce excess \mathbf{p} with a complex phase angle Ψ determined by the direction of the source $\hat{\Omega}$ and its frequency f_i :

$$\Psi_i \equiv -2\pi f_i \hat{\Omega} \cdot \Delta \vec{x} / c \quad (3)$$

Here, $\Delta \vec{x}$ is the difference in detector position and c is the speed of light. For the LIGO detectors, $|\Delta \vec{x}|/c \approx 10$ ms. Thus, a signal can, e.g., induce real negative or positive imaginary values of \mathbf{p} . Note that Ψ is time-dependent due to the rotation of the Earth. This behavior is illustrated in Figs. 1a and 1b, which show spectrograms of the real and imaginary parts of \mathbf{p} for data consisting of Advanced LIGO [36], Gaussian, Monte Carlo noise at design sensitivity plus a persistent monochromatic signal with frequency $f = 100$ Hz and strain amplitude $h_0 = 1 \times 10^{-22}$, located at (ra, dec) = (18.5 hr, +39°).

These spectrograms use $(\delta t, \Delta f) = (158 \text{ s}, 1 \text{ Hz})$ resolution and span two days of data. We utilize a coarse-graining procedure described below. The repeating pattern of light and dark at 100 Hz repeats with a period of one sidereal day. This pattern, which varies with declination and frequency, shows how the phase factor Ψ_i (Eq. 3) varies with the rotation of the Earth. It is interesting to note that these horizontal stripes are a distinct but analogous effect to the vertical stripes (also arising from radiometer phase mismatch) described in [11].

If we multiply \mathbf{p} by (an array of) the phase factor $\exp(-i\Psi_i)$ in order to “point” the radiometer map in the source direction (and reduce the injected strain amplitude to $h_0 = 4 \times 10^{-24}$) we obtain Fig. 1c. With the appropriate phase factor applied, the cross-power due to signal is real and positive. However, there is still a time-dependent modulation with a period of one sidereal day, owing to the time dependence of the detectors’ antenna response; see Eq. A2. Fig. 1d shows the signal recovered using a seedless clustering algorithm [11] to look for tracks of bright pixels. This recovery is described in greater detail in Section V.

In order to carry out a search for a transient signal associated with some direction $\hat{\Omega}$, we construct a detection statistic (see Eq. A9 in Appendix A):

$$\text{SNR}_{\text{tot}}(\hat{\Omega}) = \frac{\text{Re} \left[\sum_{i \in \Gamma} \exp(-i\Psi_i) \mathbf{p}_i \epsilon_{12}^i \right]}{\left(\sum_{i \in \Gamma} \left(\epsilon_{12}^i(\hat{\Omega}) \right)^2 \right)^{1/2}}. \quad (4)$$

The $\epsilon_{12}^i(\hat{\Omega})$ factor describes the efficiency of the detector pair [1] (see Eq. A2 in Appendix A). The sum over

i runs over some set of $(t; f)$ pixels Γ , which is determined by details of the search. A search for persistent narrowband signals, for example, might sum over time at a fixed frequency. In searches for transient signals, the set of $(t; f)$ pixels may be described, e.g., by a Bézier curve [10, 11, 37]. In the analysis that follows, we employ clustering with parametrized quadratic Bézier curves, each with a different randomly chosen direction. This family of curves includes straight lines as a subset. For fixed $\hat{\Omega}$, $\epsilon_{12}^i(\hat{\Omega})$ varies as the Earth rotates. While this effect is usually not important on timescales of $\lesssim 1000$ s, it plays a significant role for very long signals. Finally, note that Eq. 4 assumes (for the sake of compact notation) that the noise is approximately stationary and that the signal does not vary significantly with frequency; see Eq. A7. These assumptions can be straightforwardly relaxed using Eq. A6.

IV. COMPRESSION

A. Coarse-graining

In order to facilitate the search for very long-lived signals, it is desirable to compress the cross-correlation data used in long transient searches, e.g., [12]. Past analyses, targeting signals with duration ≈ 10 –1000 s, utilized (50% overlapping) 1 s-long segments. While this does not pose a significant computational challenge for the analysis of a 1500 s-long window (3000 segments), it is not practical for the analysis of 10 hr = 36000 s signals (72000 segments). First, it is currently impractical to load the data from so many segments into RAM. Second, even if we could get around the memory problem, it would take an inordinate amount of computation time to treat so many segments separately.

One solution is to average data from neighboring frequency bins in order to compress the data into a format that preserves the information needed to search for very long signals while averaging away details about short timescales. This compression, known as “coarse-graining,” enforces the assumption that the signal varies slowly enough that the change in frequency as a function of time is small compared to the bandwidth of each spectrogram pixel. Previous radiometer analyses [34, 38] have already employed a form of coarse-graining as part of a standard data conditioning procedure, but in these cases, the end data product is a single spectrum with no time-dependence (as opposed to a spectrogram), and so the coarse-graining was somewhat incidental.

We follow the coarse-graining procedure first outlined in [17] (see Eq. 5.4), which we reproduce here in slightly different notation. First, we define cross-power $C(t; f) \equiv (2/\mathcal{N}) \tilde{s}_1^*(t; f) \tilde{s}_2(t; f)$ where \mathcal{N} is a Fourier normalization constant; see Eq. A5 in Appendix A. This allows us to rewrite Eq. 2 as [47]:

$$\mathbf{p}(t; f) = \sqrt{2} C(t; f) / \sqrt{P_1'(t; f) P_2'(t; f)}. \quad (5)$$




FIG. 1: The daily modulation of \mathbf{p} . Each panel shows a $(\delta t, \Delta f) = (158 \text{ s}, 1 \text{ Hz})$ spectrogram consisting of two consecutive days of data. The data consist of Advanced LIGO [36] Monte Carlo noise plus a simulated signal at $f = 100 \text{ Hz}$. The top-left panel (a) shows $\text{Re}(\mathbf{p})$ while the top-right panel (b) shows $\text{Im}(\mathbf{p})$. In the top panels, the injected signal, for illustrative purposes, is very loud: $h_0 = 1 \times 10^{-22}$. (Since the signals are very loud, the color bar is saturated.) At $t = 0$, the phase angle is in the lower-right quadrant of the \mathbf{p} complex plane where $\text{Re}(\mathbf{p}) > 0$ and $\text{Im}(\mathbf{p}) < 0$. It rotates clockwise so that $\text{Re}(\mathbf{p})$ crosses zero and becomes negative before $\text{Im}(\mathbf{p})$ crosses zero and becomes positive. The bottom-left panel (c) shows $\text{Re}[\exp(-i\Psi)\mathbf{p}]$ for a much weaker signal: $h_0 = 4 \times 10^{-24}$. The phase factor is applied to make the signal entirely real; see Eq. A1. The bottom-right panel (d) shows the recovered signal (the most significant cluster identified in the bottom-left panel) recovered with 2×10^6 seedless clustering templates (see Section V for details). The source is located at $(\text{ra}, \text{dec}) = (18.5 \text{ hr}, +39^\circ)$ and it is recovered 7° degrees away at $(18.9 \text{ hr}, +35^\circ)$. The recovery algorithm does not assume that the signal is monochromatic.

The coarse-grained cross-power is given by:

$$C_{\text{CG}}(t; f_{\text{CG}}) \equiv \frac{\delta f}{\Delta f} \sum_{f=f_{\text{CG}}-\Delta f/2}^{f=f_{\text{CG}}+\Delta f/2} w(f) C(t; f|\hat{\Omega}). \quad (6)$$

Here, Δf is the width of the new coarse-grained bins, δf is the original fine-grained bin width, and f_{CG} are the new coarse-grained frequency bins. The variable $w(f)$ is

a weight factor, which is unity when $f \neq f_{\text{CG}} \pm \Delta f/2$. When $f = f_{\text{CG}} \pm \Delta f/2$, then $w(f)$ is some number on $[0, 1]$, which takes into account the fact that the edges of a coarse-grained bin may fall in between the edges of two fine-grained bins. The precise value is given by the fraction of the fine-grained bin that overlaps with the coarse-grained bin. Thus, e.g., a $\Delta f = 0.25 \text{ Hz}$ coarse-grained bin spanning $10.0\text{--}10.25 \text{ Hz}$ would include a $w =$

0.5 overlap with a $\delta f = 1/60$ Hz fine-grained bin spanning 9.98–10.017 Hz.

Eq. 6 can be understood as an application of Parseval's theorem. The analogous coarse-grained auto-powers, $P'_{1CG}(t; f)$, $P'_{2CG}(t; f)$ are calculated using Welch's method [39]. We thereby obtained the coarse-grained estimator:

$$p_{CG}(t; f) = \sqrt{2} C_{CG}(t; f) / \sqrt{P'_{1CG}(t; f) P'_{2CG}(t; f)}. \quad (7)$$

In addition to Eq. 6, we considered an alternative coarse-graining procedure performed in the time domain, which we refer to as “hybrid compression.” Instead of combining many narrow frequency bins associated with long segments, we combined many short segments associated with wide frequency bins. While the two procedures produced similar results, we find that the frequency-domain implementation reduced spectral leakage by up to tens of percent in power compared to the time-domain implementation, thereby allowing for better signal reconstruction. That said, there may be applications in which it is useful to carry out hybrid compression: first in the frequency domain and then in the time domain. We return to this idea at the end of the next subsection.

B. Limitations

The rotation of the Earth places a limit on the extent of possible compression [48]. If the segment duration δt is too long, the Earth will rotate significantly during the duration of the segment so that it becomes impossible to “point” the radiometer in different directions by weighting each spectrogram pixel with the appropriate phase factor; see [11]. In practice, this translates to a loss of signal-to-noise ratio.

Following [11], the fraction of signal-to-noise ratio *not* lost due to phase offset R can be written as

$$R = \cos(2\pi f \Delta \vec{x} \cdot \Delta \hat{\Omega} / c), \quad (8)$$

where f is gravitational-wave frequency, $\Delta \vec{x}$ describes the displacement between the two detectors, $\Delta \hat{\Omega}$ describes the angular error due to the rotation of the Earth, and c is the speed of light.

The angular error depends on the segment duration relative to a sidereal day. The local coordinates for directions corresponding to the poles ($\text{dec} = \pm 90^\circ$) do not change at all as the Earth rotates. The change is greatest for directions in the equatorial plane $\text{dec} = 0^\circ$, for which

$$|\Delta \hat{\Omega}| = 2\pi \frac{\delta t / 2}{24 \text{ hr}}. \quad (9)$$

Here, δt is the segment duration.

If we conservatively assume that the source is located at $\text{dec} = 0^\circ$, we can rewrite Eq. 8 as

$$R = \cos[2\pi f |\Delta \vec{x} / c| (\pi \delta t / 24 \text{ hr}) \cos(\psi)], \quad (10)$$

where ψ is the angle between $\Delta \vec{x}$ and $\Delta \hat{\Omega}$. The loss of signal-to-noise ratio is (conservatively) maximized when $\cos(\psi) = 1$. Thus, we obtain the following expression for δt :

$$\delta t = \frac{\cos^{-1}(R)}{2\pi f t_0 (\pi / 24 \text{ hr})}, \quad (11)$$

where $t_0 \equiv |\Delta \vec{x} / c|$ is the travel time between the two detectors (≈ 10 ms for the Hanford-Livingston pair). If we demand that $R > 90\%$, we obtain

$$\delta t \lesssim 197 \text{ s} \left(\frac{1000 \text{ Hz}}{f} \right) \quad (12)$$

for the Hanford-Livingston detector pair. Employing a more conservative requirement of $R > 95\%$, the prefactor becomes 140 s.

C. Cost

Now, we compare the computational cost of an analysis using uncompressed data, consisting, e.g., of 50%-overlapping $1 \text{ Hz} \times 1 \text{ s}$ spectrogram pixels, to the cost of an analysis using compressed data consisting of non-overlapping $1 \text{ Hz} \times 158 \text{ s}$ pixels. After compression, a 10 hr-long stretch of data is described by only 228 (non-overlapping) 158 s coarse-grained segments [49]. This makes it straightforward to search for \sim day-long signals. One can analyze signals with durations in excess of one week e.g., by splitting the analysis band into overlapping sub-bands.

Further, having reduced the size of the dataset by a factor of $2(158 \text{ s}) / 1 \text{ s} \approx 316$, an all-sky search using seedless clustering will run ≈ 316 times faster (than an analysis using uncompressed, 50%-overlapping, $1 \text{ Hz} \times 1 \text{ s}$ spectrograms). Seedless clustering uses phaseless templates to approximate narrowband gravitational-wave signals [10]. The sensitivity of the search, as well as its cost, is influenced by the number of templates [11].

The cost of a radiometer search is usually dominated by background estimation. The standard procedure relies on time slides, in which one data stream is repeatedly shifted in time with respect to the other in order to obtain many realizations of detector noise. Actual gravitational-wave signals do not appear in time-shifted data since the coherence is destroyed. In order to detect signals with a 3σ false alarm probability of $\approx 0.3\%$, it is necessary to analyze the ≈ 370 realizations of time-shifted data.

Scaling estimates from [11], and assuming ≈ 3 day-long spectrograms, we estimate the cost of background estimation for a realistic search to the 3σ level. We assume the search is carried out with $T_{150} = 2 \times 10^6$ templates per 150 Hz of bandwidth (a typical number of templates [11]). We find that it can be carried out using 3.6 continuously-running eight-core Intel Xeon E5-4650 CPUs or Kepler GK104 GPUs [11]. This translates to 29 continuously-

running CPU cores. The computational cost for background estimation in a realistic search scales as follows:

$$t_c \approx 10 \text{ days} \left(\frac{T_{150}}{2 \times 10^6} \right) \left(\frac{t_{\text{obs}}}{1 \text{ yr}} \right) \left(\frac{B}{1200 \text{ Hz}} \right) \left(\frac{n_{\text{ts}}}{370} \right) \left(\frac{128}{n_{\text{GPU}}} \right) \left((2)158 / \frac{\Delta f}{\delta f} \right). \quad (13)$$

Here, t_c is the computation time, t_{obs} is the duration of the observing run, B is the width of the analysis band, and n_{ts} is the number of time slides (in this case, corresponding to three-sigma confidence). The variable n_{GPU} is the number of GPUs (or, equivalently, 8-core CPUs). The ratio of $(\Delta f / \delta f)$ describes the compression achieved with coarse-graining. For the examples shown here, $(\Delta f / \delta f) = 158$. The factor of two is included if we replace overlapping windows with non-overlapping windows.

D. Hybrid compression

At the end of the previous subsection, we noted that there may be applications for which it is helpful to carry out hybrid compression: first in the frequency domain and then in the time domain. For example, if we are interested a particular direction in the sky, one can combine many frequency-domain coarse-grained segments into even longer segments via Eq. 4. Since, by assumption, each of the frequency-domain coarse-grained segments in the sum already takes into account the phase offset between the two detectors, the constraint from Eq. 12 does not apply, and so there is no limit to the number of segments that can be combined. In this way, it should be possible to represent an arbitrarily long span of data with a fixed number of spectrogram pixels, facilitating a search for a signal of any duration. It is important to note, however, that this hybrid technique can only be applied by assuming a specific source location, and so it does not lend itself to an efficient all-sky search.

E. Trade-offs

By coarse-graining spectrographic data with Eq. 6, we are trading reduced frequency resolution for reduced computational burden. As a consequence, the sensitivity of the search is reduced in comparison to an idealized lossless search using arbitrarily large computing resources. To see why, consider a narrowband signal passing through just one fine-grained spectrogram pixel (width = δf) during each segment in time. When the data are coarse-grained to a width Δf , the signal in this single fine-grained bin is averaged with data in neighboring frequency bins containing only noise. In this worst-case scenario, the detection statistic SNR_{tot} is reduced by a factor $\sqrt{\Delta f / \delta f}$ —corresponding to a factor of ≈ 13 for the $1 \text{ Hz} \times 158 \text{ s}$ pixels used for our demonstration below.

In practice, the loss of sensitivity can be less if the signal is not confined to a single fine-grained bin at any given time. Also, fine-grained bins are sometimes contaminated with noise artifacts including non-Gaussian and non-stationary noise, which are, to some extent, ameliorated through coarse-graining. Such noise artifacts can therefore serve to make the loss of sensitivity less severe in comparison to the idealized case.

Once the data are coarse-grained, the high-resolution information is lost; the process is not invertible. However, in the event of a candidate detection, one can imagine reanalyzing the original lossless dataset with a targeted search, looking only in a small band and in a specific direction identified by the coarse-grained search. By severely restricting the parameter space of the search, it is possible to carry out a computationally feasible follow-up study using fine-grained data with reasonable computational resources.

V. DEMONSTRATION

Our goal here is to show that very long signals, lasting hours or longer, can be successfully identified by employing an existing pattern recognition algorithm on compressed data. For our first demonstration, we recover the signal shown in Fig. 1c. We create a simulated dataset spanning two days of Monte Carlo data for the Advanced LIGO Hanford and LIGO Livingston detectors operating at design sensitivity. The data consists of Gaussian noise plus a simulated monochromatic signal with frequency $f = 100 \text{ Hz}$ and strain amplitude $h_0 = 4 \times 10^{-24}$, located at $(\text{ra}, \text{dec}) = (18.5 \text{ hr}, +39^\circ)$. The signal is circularly polarized.

The data are cross-correlated to produce spectrograms with resolution of $(\delta t, \Delta f) = (158 \text{ s}, 1 \text{ Hz})$ with non-overlapping segments. We consider an observing band spanning 10–300 Hz. This spectrogram (with a phase factor applied to point the radiometer in the appropriate direction) is shown in Fig. 1c. Once the data are compressed, we carry out a seedless clustering search [10, 11].

We make no assumptions about the start time, stop time, frequency, or sky location of the search—subject to the constraint that the signal morphology can be approximated as a cubic Bézier curve of at least 1.8 hr in duration. (Following [11], we take advantage of the fact that \mathbf{p} encodes directional information as a complex phase in order to point each seedless template in a different random direction.) We record SNR_{tot} for 2×10^6 randomly-generated templates in order to identify the template that best fits the signal. The best-fit template (with $\text{SNR}_{\text{tot}} = 56$) is shown in Fig. 1d. The recovery represents a reasonably good match to the injected signal. (It would be better recovered using a larger number of templates.) Running on an eight-core CPU, this 48 hr stretch of data was analyzed in $\approx 550 \text{ s}$.

As a second test, we attempted to recover a signal with a more complicated time-frequency evolution. The signal

turns on ≈ 5 hr after the beginning of data-taking and ends near $t \approx 17$ hr during which time the signal evolves according to:

$$\begin{aligned}\Phi(t) &= \Phi_0 + 2\pi \left[f_0 (t - t_0) + \frac{1}{2} \dot{f} (t - t_0)^2 \right] \\ h_+ &= h_0 \cos(\Phi(t)) \\ h_\times &= h_0 \sin(\Phi(t)).\end{aligned}\quad (14)$$

We choose $f_0 = 100$ Hz, $\dot{f} = 20$ Hz day $^{-1}$, and $h_0 = 4 \times 10^{-24}$. The signal is added to simulated noise and processed using the same procedure used for the first test injection.

In Fig. 2a, we show a spectrogram of $\rho(t; f)$. The signal can be seen as a faint track beginning at 5 hr and appearing to end at 12 hr. (The last several hours of the signal are difficult to see by eye since the alignment of the detectors makes it weak during this span.) The recovered signal is shown in Fig. 2b. Once again, we use 2×10^6 cubic Bézier templates. As before, the source is located at (ra, dec) = (18.5 hr, +39°), though, this information is not provided to the search algorithm. The recovered track (with $\text{SNR}_{\text{tot}} = 11$) is a good match.

In Fig. 2c, we show the spectrogram for an even longer week-long signal present in two weeks of data. The signal parameters are identical to parameters used in Figs. 2ab except we use a smaller value of $\dot{f} = 10$ Hz day $^{-1}$. Using 2×10^6 cubic Bézier templates, we recover the injected signal in Fig. 2d with $\text{SNR}_{\text{tot}} = 67$.

In Table I, we show the strain values h_0 for which different signals (such as those shown in Fig. 2) can be detected in 90 days of data with a false alarm probability of 1% and a false dismissal probability of 50%. We look for signals in spectrograms of duration $\Delta t = 1$ day, and in one case, $\Delta t = 14$ days. We simulate noise to determine the threshold of SNR_{tot} required for 1% false-alarm probability detection. Then, we recover simulated signals with varied h_0 to determine the value for which at least half of the signals are recovered with a sufficiently large value of SNR_{tot} (false dismissal probability = 50%). In all cases, we assume an optimally oriented source with an optimal sky location. (For randomly oriented sources with random sky locations, we expect h_0 to increase by a factor of ≈ 1.5 –2; see, e.g., [11].) We employ 2×10^5 templates in a 300 Hz bandwidth.

The results shown in Table I can be understood in terms of a simple scaling law, which can be used to guide our expectations for the detectability of an arbitrary signal. If the seedless clustering template approximately matches the signal we are trying to recover, then the detection statistic scales approximately as follows:

$$\text{SNR}_{\text{tot}} \propto \frac{h_0^2}{\overline{P(f)}} \left(\frac{t_{\text{dur}}}{\Delta f} \right)^{1/2}, \quad (15)$$

where $\overline{P(f)}$ is the strain noise power spectral density averaged over the emission frequency and t_{dur} is the duration of emission. While t_{dur} , Δf , and $\overline{P(f)}$ are the

f_0 (Hz)	\dot{f} (Hz day $^{-1}$)	t_{dur} (days)	Δt (days)	h_0
100	20	0.29	1	3.6×10^{-24}
100	0	0.29	1	3.3×10^{-24}
100	10	7	14	1.7×10^{-24}
1000	0	0.29	1	6.2×10^{-24}

TABLE I: The detectability of different signals in simulated Advanced LIGO noise. The variables f_0 and \dot{f} describe the signal morphology following Eq. 14 while t_{dur} is the signal duration. Δt is the duration of each spectrogram input to the seedless clustering algorithm. The variable h_0 is the signal strength that can be recovered with a false alarm probability of 1% and a false dismissal probability of 50% in 90 days of data. We assume optimally oriented sources in optimal sky locations. We employ 2×10^5 templates per spectrogram and analyze data in a 300 Hz bandwidth. We use $(\delta t, \Delta f) = (158 \text{ s}, 1 \text{ Hz})$ pixels. The sensitivities quoted here can very likely be improved by using more templates.

most important factors in determining the detectability of a signal, other variables play a role too. For example, using a fixed number of templates, it is harder to detect signals as the frequency increases (all else equal) because the shrinking diffraction-limited resolution makes it more challenging to match the signal to a template [11].

Together, the recoveries demonstrated in Figs. 1 and 2, and also Tab. I, show that it is possible to recover very long-lived signals with a range of different signal morphologies in compressed cross-correlated data. The recovery works for relatively weak signals and can be carried out with reasonable computational resources. It is made feasible by coarse-graining, which allows us to analyze days of data using a single spectrogram.

In order to place the results of Tab. I in context, we consider an illustrative emission scenario. We consider a neutron star in a binary, which emits gravitational waves, for example, through r -mode instabilities, which decay on a timescale of a ≈ 1 week. By assuming accretion-torque balance, it has been hypothesized [31] that the low-mass X-ray binary Scorpius X-1 emits gravitational waves with a strain of $\approx 9 \times 10^{-26}$ at $f = 100$ Hz [40]. By tuning the bin width to match the observed Doppler modulation of Scorpius X-1, we expect the sensitivity stated in Tab. I could be improved by a factor of $(1 \text{ Hz}/0.05 \text{ Hz})^{1/4} \approx 2$ near 100 Hz [50]. This is still about a factor of 10 below the expected strain from Scorpius X-1 assuming accretion-torque balance.

Thus, assuming the hypothesized accretion-torque balance scenario, advanced detectors, using the method presented here, are unlikely to detect very long transient gravitational waves from a source like Scorpius X-1 ($d = 2.3$ kpc). Detection would require a source that is either significantly closer or significantly brighter. That said, it is important to note that detecting *any* gravitational waves from neutron stars such as Scorpius X-1 is a challenging proposition for a variety of search techniques [41]. The method proposed here will help to cover the full parameter space of possible signals from such




FIG. 2: Recovery of a very long transient signal with seedless clustering. Top-left (a): a $(\delta t, \Delta f) = (158 \text{ s}, 1 \text{ Hz})$ spectrogram of $\rho(t; f|\hat{\Omega})$ (Eq. A3) consisting of one day of data. The data consist of Advanced LIGO Monte Carlo noise plus a simulated signal. The signal begins at $f = 100 \text{ Hz}$ and increases with a spin-up given by $\dot{f} = 20 \text{ Hz day}^{-1}$. The amplitude is $h_0 = 4 \times 10^{-24}$. The signal persists between $t \approx 5$ –17 hr. (It grows very faint after $t \approx 12 \text{ hr}$ due the unfavorable orientation of the Earth during this period.) Top-right (b): the recovered signal recovered with 2×10^6 cubic Bézier templates; $\text{SNR}_{\text{tot}} = 11$. Bottom-left (c): a $(\delta t, \Delta f) = (158 \text{ s}, 1 \text{ Hz})$ spectrogram of $\rho(t; f|\hat{\Omega})$ consisting of two weeks of data. The data consist of Advanced LIGO Monte Carlo noise plus a simulated signal. The signal begins at $f = 100 \text{ Hz}$ and increases with a spin-up given by $\dot{f} = 10 \text{ Hz day}^{-1}$. The amplitude is $h_0 = 4 \times 10^{-24}$. The signal persists between $t \approx 3$ –10 days. Bottom-right (d): the recovered signal the most significant cluster identified in the bottom-left panel) recovered with 2×10^6 cubic Bézier templates; $\text{SNR}_{\text{tot}} = 67$.

sources.

Before we move on, we note in passing, an interesting difference between this very long transient algorithm and the matched filtering algorithm proposed in [13]. In [13], the authors note that—at a fixed frequency—the *matched filtering* signal-to-noise ratio depends only on the total energy emitted in gravitational waves E_{gw} and not on the signal duration t_{dur} . However, since radio-

metric signal-to-noise ratio scales like $\text{SNR} \propto \sqrt{t_{\text{dur}}} h_0^2$, and since $E_{\text{gw}} \propto t_{\text{dur}} h_0^2$, $\text{SNR} \propto E_{\text{gw}}/\sqrt{t_{\text{dur}}}$. Thus, for radiometer analyses, we obtain a different scaling behavior than the matched filter behavior noted in [13]: given a fixed energy budget, shorter gravitational-wave transients are easier to detect than longer transients. Given a constant flux $\Phi_{\text{gw}} \propto dE_{\text{gw}}/dt$, the signal-to-noise ratio grows with time: $\text{SNR} \propto \Phi_{\text{gw}} \sqrt{t_{\text{dur}}}$.

VI. CONCLUSIONS

We have demonstrated the machinery of an efficient search for very long lived transient gravitational waves. Such a search will be useful, both looking for very long transient signals from neutron stars, and for understanding the time-dependence of signals identified through persistent radiometer searches and stochastic background searches. Our procedure relies on coarse-graining to create an intermediate data product that is computationally manageable to work with. In carrying out our demonstration, we explored the modulations in cross-correlated data due to the rotation of the Earth. Future work, building on this demonstration, will explore new physics (and instrumental effects) on timescales of hours to days.

We note that the method described in [13], which builds on existing searches for continuous waves from persistent emitters, can likely achieve better sensitivity than the method proposed here in its domain of utility: sources following the canonical spin evolution typical of an isolated pulsar. This is because continuous wave searches for isolated neutron stars (see, e.g., [42]) assume a heavily constrained parameter space of pulsar-like signals. However, there are important differences between radiometry and continuous waves searches, which make our new method highly complementary.

Like all applications of gravitational-wave radiometry, the method described here is most useful in cases when the signal does not match a simple template. For neutron stars in binary systems, radiometer measurements [34, 43] can be comparable in sensitivity to continuous wave techniques [44, 45], which rely on a more detailed signal model. Further, radiometer searches, including the one described here, are extremely robust because they do not make assumptions about the phase evolution of the signal. Finally, in the method described here, we allow for significant variation in the frequency evolution of the gravitational-wave signal, which allows us to probe a completely different parameter space compared to continuous wave searches, which assume very limited changes in the neutron star spin: $\lesssim 6 \times 10^{-9} \text{ Hz s}^{-1}$ versus $\lesssim 300 \text{ Hz s}^{-1}$.

While we have focused here on signals lasting hours to weeks, this method can be extended to target both shorter and longer signals as well. While significant work has been carried out to develop tools for probing signals with durations of 10–1000 s, it may prove useful to apply some modest coarse-graining ($\Delta f / \delta f \gtrsim 1$) to searches that are stretched for computational resources. Characterization of \approx year-long signals may require some additional development since a year-long spectrogram of compressed data still requires significant memory resources. However, extrapolating from our analysis here of a two-week-long spectrogram, it seems likely that such a search could be carried out through a number of options such as the use of computers with larger amounts of memory $\gtrsim 24 \text{ GB}$ and/or working with narrower frequency bands.

Another next step is to consider how realistic searches,

using the ideas presented here, can be optimized. For example, a search for persistent narrowband signals may benefit from the use of overlapping frequency bins and/or frequency bins of different sizes. Such optimizations would be equally useful for any radiometer search.

VII. ACKNOWLEDGMENTS

We thank Joe Romano for helpful discussion on coarse-graining, Vladimir Dergachev for discussion on intermittent pulsars, and David Keitel for carefully reading and commenting on an earlier version of the manuscript. ET was a member of the LIGO Laboratory, supported by funding from United States National Science Foundation. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. VM's work was supported by NSF grant PHY1204944. NC's work was supported by NSF grant PHY-1204371. This is LIGO document P1400256.

Appendix A: Details of radiometry

1. Definitions

Given two simultaneous strain measurements s_1 and s_2 , the cross-correlation statistic \hat{Y} (and its associated uncertainty σ) can be written as [1]:

$$\begin{aligned}\hat{Y}(t; f | \hat{\Omega}) &= \frac{2}{\mathcal{N}} \text{Re} \left[e^{2\pi i f \hat{\Omega} \cdot \Delta \vec{x} / c} \frac{1}{\epsilon_{12}(t | \hat{\Omega})} \tilde{s}_1^*(t; f) \tilde{s}_2(t; f) \right] \\ \sigma^2(t; f | \hat{\Omega}) &= \frac{1}{2} \left| \frac{1}{\epsilon_{12}(t | \hat{\Omega})} \right|^2 P_1'(t; f) P_2'(t; f).\end{aligned}\tag{A1}$$

Here, $\tilde{s}_I(t; f)$ is the discrete Fourier transform of the strain series from detector I and $P'(t; f)$ is the auto-power spectrum calculated using neighboring data segments. Each segment has duration δt . The argument t is the starting time of the data segment and f is frequency. \mathcal{N} is a Fourier normalization coefficient, $\hat{\Omega}$ is the direction of the source, $\Delta \vec{x}$ is the difference in detector position, and c is the speed of light. Above, we carry out simulations for the Advanced LIGO network (consisting of detectors in Hanford, WA and Livingston, LA) for which $|\Delta \vec{x}| / c \approx 10 \text{ ms}$. The $\epsilon_{12}(t | \hat{\Omega})$ factor describes the efficiency of the detector pair [1]:

$$\epsilon_{12}(t | \hat{\Omega}) \equiv \frac{1}{2} \sum_A F_1^A(t | \hat{\Omega}) F_2^A(t | \hat{\Omega}).\tag{A2}$$

Here, $F_I^A(t | \hat{\Omega})$ is the antenna factor [46] for detector I and $A = +, \times$ are the different polarization states.

It is useful to define “pixel signal-to-noise ratio”:

$$\rho(t; f|\hat{\Omega}) = \hat{Y}(t; f|\hat{\Omega}) / \sigma(t; f|\hat{\Omega}). \quad (\text{A3})$$

We use $\rho(t; f|\hat{\Omega})$ for its convenience as a tool for visual representation of cross-correlated data and also because it allows for compact derivations. In the context of all-sky searches, it is also useful to define “complex signal-to-noise ratio” [11]

$$\mathbf{p}(t; f) = \kappa \tilde{s}_1^*(t; f) \tilde{s}_2(t; f) / \sqrt{P_1'(t; f) P_2'(t; f)}, \quad (\text{A4})$$

where, for the sake of compactness, we define a normalization factor

$$\kappa \equiv 2\sqrt{2}/\mathcal{N}. \quad (\text{A5})$$

Comparing Eq. A3 and A4, we see that $\mathbf{p}(t; f)$ is the same as $\rho(t; f)$ except no phase factor has been applied, and we do not take the real part. In doing so, we obtain a data product with information from all directions on the sky.

2. Clustering

Having defined spectrograms of ρ or \mathbf{p} (Eqs. A3 and A4), the next step in a radiometer analysis is to look for clusters of statistically significant excess coherence. Details of the clustering procedure depend on the specific search, but there are certain common features.

First, a rule is defined for how clusters can be formed. In the case of the targeted narrowband radiometer [16, 34, 38], the signal is assumed to persist with a fixed frequency, and so each cluster corresponds to the set of all pixels at a given frequency; i.e., one spectrogram row.

Once a cluster is defined, all the pixels in that cluster are combined in order to produce a detection statistic for that cluster denoted SNR_{tot} . The detection statistic (for a targeted search with a known sky location) can be written as a weighted average [16, 34, 38]:

$$\text{SNR}_{\text{tot}} = \frac{\sum_{i \in \Gamma} \hat{Y}_i \sigma_i^{-2}}{(\sum_{i \in \Gamma} \sigma_i^{-2})^{1/2}}. \quad (\text{A6})$$

(We have suppressed the dependence on $\hat{\Omega}$ for the sake of brevity.) The sum over i runs over some set of $(t; f)$ pixels Γ , which is determined by details of the search. A search for persistent narrowband signals, for example, might sum over time at a fixed frequency. In searches for transient signals, the set of $(t; f)$ pixels may be described, e.g., by a Bézier curve [10, 11, 37].

If we assume for the sake of simplicity that the noise is approximately stationary, and that the signal does not

vary significantly in frequency, then we can write

$$\begin{aligned} \hat{Y}_i &= \rho_i \sigma_0 / \epsilon_{12}^i \\ \sigma_i &= \sigma_0 / \epsilon_{12}^i. \end{aligned} \quad (\text{A7})$$

Recall that ϵ_{12}^i (Eq. A2) is an efficiency factor arising from the detector response function, which changes in time as the Earth rotates. The variable σ_0 is some number determined only by the detector strain sensitivities. Combining Eqs. A6 and A7, we obtain:

$$\text{SNR}_{\text{tot}} = \frac{\sum_{i \in \Gamma} \rho_i \epsilon_{12}^i}{(\sum_{i \in \Gamma} (\epsilon_{12}^i)^2)^{1/2}}. \quad (\text{A8})$$

3. All-sky search

To this point in the appendix, we have assumed that the sky location $\hat{\Omega}$ is known a priori. This assumption holds for targeted searches, e.g., searching for gravitational waves from the low-mass X-ray binary, Scorpius X-1 [31]. When the sky location is not known, the search becomes more complicated. To see this, consider an all-sky search which considers some finite set of directions $\{\hat{\Omega}_i\}$. If we look in some direction $\hat{\Omega}_i$, which is offset from the true source direction $\hat{\Omega}_t$, then the phase factor in Eq. A1 will not rotate the signal power perfectly into the real number axis. As a result, signal power leaks into the imaginary direction, and can even produce negative power on the real axis.

In order to carry out an efficient all-sky search, we can use Eqs. A8 and A4 to define a detection statistic that can quickly scan over sky locations:

$$\text{SNR}_{\text{tot}}(\hat{\Omega}) = \frac{\text{Re} [\sum_{i \in \Gamma} \exp(-i\Psi_i) \mathbf{p}_i \epsilon_{12}^i]}{(\sum_{i \in \Gamma} (\epsilon_{12}^i)^2)^{1/2}}. \quad (\text{A9})$$

(This definition of SNR_{tot} is essentially the same one that is proposed in [11].) Graphically, Eq. A9 can be interpreted as applying an $\hat{\Omega}$ -dependent phase array to a complex-valued spectrogram (as in Figs. 1a and 1b) before summing over clusters of pixels.

At first glance, Eq. A9 may appear tautological. There is already implicit dependence on $\hat{\Omega}$ in Eq. A8 hiding in ρ_i , ϵ_i , and therefore in SNR_{tot} . However, the ingredients used in Eq. A8 are calculated for a specific sky location $\hat{\Omega}$ in such a way that the information about all other sky locations is lost when the imaginary part is discarded in Eq. A1. Thus, the formulation described by Eq. A9 allows us to search the entire sky with a single data product. Using parallel computing architecture, we can efficiently search many directions in parallel with a single manageable array, dramatically speeding up calculations in the process [11].

-
- [1] E. Thrane, S. Kandhasamy, C. D. Ott, et al., Phys. Rev. D **83**, 083004 (2011).
 - [2] A. L. Piro and E. Pfahl, Astrophys. J. **658**, 1173 (2007).
 - [3] M. H. P. M. van Putten, Astrophys. J. Lett. **684**, L91 (2008).
 - [4] K. Kiuchi, M. Shibata, P. J. Montero, and J. A. Font, Phys. Rev. Lett. **106**, 251102 (2011).
 - [5] P. Raffai, Z. Frei, Z. Márka, and S. Márka, Classical Quantum Gravity **24**, S457 (2007).
 - [6] D. Murphy, M. Tse, P. Raffai, I. Bartos, R. Khan, Z. Márka, L. Matone, K. Redwine, and S. Márka, Phys. Rev. D **87**, 103008 (2013).
 - [7] A. Corsi and P. Mészáros, Astrophys. J. **702**, 1171 (2009).
 - [8] A. L. Piro and E. Thrane, Astrophys. J. **761**, 63 (2012).
 - [9] A. L. Piro and C. D. Ott, Astrophys. J. **736**, 108 (2011).
 - [10] E. Thrane and M. Coughlin, Phys. Rev. D **88**, 083010 (2013).
 - [11] E. Thrane and M. Coughlin, Phys. Rev. D **89**, 063012 (2014).
 - [12] J. Aasi et al., Phys. Rev. D **88**, 122004 (2013).
 - [13] R. Prix, S. Giampanis, and C. Messenger, Phys. Rev. D **84**, 023007 (2011).
 - [14] A. Arvanitaki, M. Baryakhtar, and X. Huang, Phys. Rev. D **91**, 084011 (2015).
 - [15] J. Aasi et al., Phys. Rev. Lett. **113**, 231101 (2014).
 - [16] S. W. Ballmer, Classical Quantum Gravity **23**, S179 (2006).
 - [17] B. Abbott et al., Phys. Rev. D **69**, 122004 (2004).
 - [18] C. A. van Eysden and A. Melatos, Classical Quantum Gravity **25**, 225020 (2008).
 - [19] M. F. Bennett, C. A. van Eysden, and A. Melatos, MNRAS **409**, 1705 (2010).
 - [20] A. G. Lyne, S. L. Shemar, and F. G. Smith, MNRAS **315**, 534 (2000).
 - [21] K. Kashiyaama and K. Ioka, Phys. Rev. D **83**, 081302 (2011).
 - [22] Y. Levin and M. van Hoven, MNRAS **418**, 659 (2011).
 - [23] P. Jaranowski, A. Królak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998).
 - [24] D. I. Jones and N. Andersson, MNRAS **331**, 203 (2002).
 - [25] L. Gualtieri, R. Ciolfi, and V. Ferrari, *Structure, deformations and gravitational wave emission of magnetars* (2011).
 - [26] S. Dall’Osso, S. N. Shore, and L. Stella, MNRAS **398**, 1869 (2009).
 - [27] N. Andersson and K. D. Kokkotas, Int. J. Mod. Phys. D **10**, 381 (2001).
 - [28] R. Bondarescu, S. A. Teukolsky, and I. Wasserman, Phys. Rev. D **79**, 104003 (2009).
 - [29] A. Melatos and M. Priymak, Astrophys. J. **794**, 170 (2014).
 - [30] F. Camilo, S. M. Ransom, S. Chatterjee, S. Johnston, and P. Demorest, Astrophys. J. **746**, 63 (2012).
 - [31] L. Bildsten, Astrophys. J. Lett. **501**, L89 (1998).
 - [32] E. Thrane, N. Christensen, and R. M. S. Schofield, Phys. Rev. D **87**, 123009 (2013).
 - [33] E. Thrane, N. Christensen, R. M. S. Schofield, and A. Effler, Phys. Rev. D **90**, 023013 (2014).
 - [34] B. P. Abbott (The LIGO Scientific and Virgo Collaborations), Phys. Rev. Lett. **107**, 271102 (2011).
 - [35] B. Abbott et al., Nature **460**, 990 (2009).
 - [36] Harry, G. M. for the LIGO Scientific Collaboration, Classical Quantum Gravity **27**, 084006 (2010).
 - [37] M. Coughlin, E. Thrane, and N. Christensen, Phys. Rev. D **90**, 083005 (2014).
 - [38] B. Abbott et al., Phys. Rev. D **76**, 082003 (2007).
 - [39] P. D. Welch, IEEE Transactions on Audio Electroacoustics **AU-15**, 70 (1967).
 - [40] L. Sammut, C. Messenger, A. Melatos, and B. J. Owen, Phys. Rev. D **89**, 043001 (2014).
 - [41] C. Messenger et al. (2015), <http://arxiv.org/abs/1504.05889>.
 - [42] J. Abadie et al., Phys. Rev. D **85**, 022001 (2012).
 - [43] E. Thrane, S. Mitra, N. Christensen, V. Mandic, and A. Ain (2015), <http://arxiv.org/abs/1504.02158>.
 - [44] Aasi et al., Phys. Rev. D **90**, 062010 (2014).
 - [45] J. Aasi et al. (The LIGO Scientific and Virgo Collaborations), Phys. Rev. D **91**, 062008 (2015).
 - [46] K. S. Thorne, in *Three Hundred Years of Gravitation*, edited by S. W. Hawking and W. Israel (Cambridge University Press, Cambridge, 1989).
 - [47] In our calculation of discrete Fourier transforms, we employ a Tukey window with a 0.5 s rise and fall time.
 - [48] An independent, but weaker constraint arises by taking into account Doppler modulation. In the case of a half-year-long signal (for which Doppler modulation is maximal), the modulation depth is $\approx 10^{-4} f$. Thus, a radiometer search for months-long signals at $f \approx 1000$ Hz should employ a frequency resolution $\Delta f \gtrsim 0.1$ Hz.
 - [49] Overlapping segments are often used in long-transient searches to recover signal power lost to Hann windowing, which, in turn, is carried out to reduce spectral leakage. Here, we have replaced Hann windows with a rapidly rising Tukey window. This allows us to minimize spectral leakage without a significant reduction in signal power, thereby facilitating the use of non-overlapping segments.
 - [50] Here we are envisioning a targeted search, which, if necessary, could employ hybrid compression.